Reliable and efficient solution of genome-scale models of Metabolism and macromolecular Expression

Ding Ma*, Laurence Yang†, Ronan M. T. Fleming‡, Ines Thiele‡,
Bernhard O. Palsson†§ & Michael A. Saunders*

September 28, 2016

arXiv:1606.00054v2 [q-bio.MN] 26 Sep 2016

Constraint-Based Reconstruction and Analysis (COBRA) is currently the only methodology that permits integrated modeling of Metabolism and macromolecular Expression (ME) at genome-scale. Linear optimization computes steady-state flux solutions to ME models, but flux values are spread over many orders of magnitude. Standard double-precision solvers may return inaccurate solutions or report that no solution exists. Exact simplex solvers are extremely slow and hence not practical for ME models that currently have 70,000 constraints and variables and will grow larger. We have developed a quadruple-precision version of our linear and nonlinear optimizer MINOS, and a solution procedure (DQQ) involving Double and Quad MINOS that achieves efficiency and reliability for ME models. DQQ enables extensive use of large, multiscale, linear and nonlinear models in systems biology and many other applications.

Constraint-Based Reconstruction and Analysis (COBRA) [1] has been applied successfully to predict phenotypes for a range of genome-scale biochemical processes. The popularity of COBRA is partly due to the efficiency of the underlying optimization algorithms, permitting genome-scale modeling at a particular timescale using readily available open source software [2, 3] and industrial quality optimization algorithms [19, 5, 6]. A widespread application of COBRA is the modeling of steady states in genome-scale Metabolic models (M models). COBRA has also been used to model steady states in macromolecular Expression networks (E models), which stoichiometrically represent the transcription, translation, post-translational modification and formation of all protein complexes required for macromolecular biosynthesis and metabolic reaction catalysis [7, 8]. COBRA of metabolic networks or expression networks depends on numerical optimization algorithms to compute solutions to certain model equations, or to determine that no solution exists. Our purpose is to discuss available options and to demonstrate an approach that is reliable and efficient for ever larger networks.

Metabolism and macromolecular Expression (ME) models have opened a whole new vista for predictive mechanistic modeling of cellular processes, but their size and multiscale nature pose a challenge to standard linear optimization (LO) solvers based on 16-digit double-precision floating-point arithmetic. Standard LO solvers usually apply scaling techniques [15, 10] to problems that are not already well scaled. The scaled problem typically solves more efficiently and accurately, but the solver must then unscale the solution, and this may generate significant primal or dual infeasibilities in the original problem (the constraints or optimality conditions may not be accurately satisfied).

A *lifting* approach [21] has been implemented to alleviate this difficulty with multiscale problems. Lifting reduces the largest matrix entries by introducing auxiliary constraints and variables. This approach has permitted standard (double-precision) LO solvers to find more accurate solutions, even though the final objective value is still not satisfactory (Table 3). Another approach to increasing the precision is to use an exact solver. An exact simplex solver QSopt_ex [12, 9] has been used for a ME model of *Thermotoga maritima* [4] (model TMA_ME) representing a network with about 18,000 metabolites and reactions. The solution time was about two weeks, compared to a few minutes (Table 2) for a standard double-precision solver, but the latter's final objective value had only one correct digit (Table 3). QSopt_ex has since been applied to a collection of 98 metabolic models by Chindelvitch et al. [15] via their MONGOOSE toolbox. Most of the 98 models have less than 1000 metabolites and reactions. QSopt_ex required about a day to solve all models [15], compared to a few seconds in total for a standard solver.

To advance COBRA for increasingly large biochemical networks, solvers that perform more efficiently than exact solvers and also perform more reliably than standard LO solvers are definitely needed. Gleixner et al.

*Dept of Management Science and Engineering, Stanford University, Stanford, CA, USA. Correspondence should be addressed to M. S. (saunders@stanford.edu).

†Dept of Bioengineering, University of California, San Diego, CA, USA.

‡Luxembourg Centre for Systems Biomedicine, University of Luxembourg, Luxembourg.

§Novo Nordisk Foundation Center for Biosustainability, The Technical University of Denmark, 2970 Hørsholm, Denmark.

[16, 17] have addressed this need, and Chapter 4 of [17] is devoted to multiscale metabolic networks, showing significant improvement relative to CPLEX [5]. Our work is complementary and confirms the value of enhancing the simplex solver in [16, 17] to employ quadruple-precision computation, as we have done here.

Let Single/Double/Quad denote the main floating-point options, with about 7, 16, and 34 digits of precision respectively. For many years, scientific computation has advanced in two complementary ways: improved *algorithms* and improved *hardware*. Compilers have typically evaluated expressions using the same arithmetic as the variables' data type. Most scientific codes apply Double variables and Double arithmetic throughout (16 significant digits stored in 64-bit words). The floating-point hardware often has *slightly extended precision* (80-bit registers). Kahan [18] notes that early C compilers generated Double instructions for all floating-point computation *even for program variables stored in single precision*. Thus for a brief period, C programs were serendipitously more reliable than typical Fortran programs of the time. (For Single variables a and b, Fortran compilers would use Single arithmetic to evaluate the basic expressions $a \pm b$, $a*b$, a/b, whereas C compilers would transfer a and b to longer registers and operate on them using Double arithmetic.) Most often, the C compiler's extra precision was not needed, but occasionally it did make a critical difference. Kahan calls this the *humane* approach to debugging complex numerical software. Unfortunately, Quad hardware remains very rare and for the foreseeable future will be simulated on most machines by much slower software. Nevertheless, we believe the time has come to produce Quad versions of key sparse-matrix packages and large-scale optimization solvers for multiscale problems.

Here, we report the development and biological application of Quad MINOS, a quadruple-precision version of our general-purpose, industrial-strength linear and nonlinear optimization solver MINOS [10, 11]. We also developed a Double-Quad-Quad MINOS procedure (DQQ) that combines the use of Double and Quad solvers in order to achieve a balance between efficiency in computation and accuracy of the solution. We extensively tested this DQQ procedure on 83 genome-scale metabolic network models (M models) obtained from the UCSD Systems Biology repository [21, 22] and 78 from the BiGG database [1]. We also applied DQQ to ME models of *Thermotoga maritima* [4] (about 18,000 metabolites and reactions) and *E. coli* K12 MG1655 [24] (about 70,000 metabolites and reactions). For M models, we find that Double MINOS alone is sufficient to obtain non-zero steady-state solutions that satisfy feasiblility and optimality conditions with a tolerance of 10^{-7}. For ME models, application of our DQQ procedure resulted in non-zero steady-state solutions that satisfy feasibility and optimality conditions with a tolerance of 10^{-20}. The largest model, a lifted version of the *E. coli* ME model, required 4.5 hours, while an exact solver would take months.

Thus, we expect our DQQ procedure to be a robust and efficient tool for the increasingly detailed study of biological processes, such as metabolism and macromolecular synthesis, and many other scientific fields.

Results

Efficient combination of Double and Quad. To achieve reliability and efficiency on multiscale problems, in general form as *linear optimization* (LO):

$$\min\ c^T v \ \text{ s.t. } \ Sv = 0, \ \ \ell \le v \le u, \tag{1}$$

we developed the following 3-step procedure.

DQQ procedure.

Step D (Cold start in Double with scaling): Apply Double MINOS with moderately strict options. Save a final basis file.

Step Q1 (Warm start in Quad with scaling): Start Quad MINOS from the saved file with stricter options. Save a final basis file.

Step Q2 (Warm start in Quad without scaling): Start Quad MINOS from the second saved file with no scaling but stricter LU options.

DQQ is described further in Algorithm 1, where loop **1** is the primal simplex method, P is a permutation matrix, and δ_1, δ_2 are Feasibility and Optimality tolerances. MINOS terminates loop **1** when the (possibly scaled) bounds on v are satisfied to within δ_1 and $z_j/(1+\|y\|_\infty)$ has the correct sign to within δ_2. Table 1 shows the default runtime options for Double MINOS and the options chosen for each step of DQQ. Scale specifies whether the problem data should be scaled before the problem is solved (and unscaled after). Tolerances δ_1, δ_2 specify how well the primal and dual constraints of the (possibly scaled) problem should be satisfied. Expand frequency controls the MINOS anti-degeneracy procedure [16]. The LU tolerances balance stability and sparsity when LU factors of B are computed or updated.

Steps D and Q1 are usually sufficient, but if Q1 is interrupted, Q2 provides some insurance and ensures that the tolerances δ_1 and δ_2 are imposed upon the original problem (not the scaled problem). For conventional Double solvers, it is reasonable to set tolerances in the range 10^{-6} to 10^{-8}. For Quad MINOS, we set $\delta_1 = \delta_2 = 10^{-15}$ to be sure of capturing reaction fluxes v_j as small as $O(10^{-10})$.

Small M models. Of the 98 metabolic network models in the UCSD Systems Biology repository [21], A. Ebrahim was able to parse 83 models [26] and compute solutions with a range of solvers [27]. We constructed MPS files for the 83 models [22] and solved them via DQQ. Most models have less than 1000 metabolites and reactions. Almost all models solved in less than 0.08 seconds, and many in less than 0.01 seconds. The total time was less than 3 seconds. In contrast, the exact arithmetic solver needs a day [15].

Algorithm 1: DQQ procedure

Data: Linear program (1)

Result: Flux vector v^*, basis partition $SP=(B\ N)$, one of three states: optimal, infeasible, or unbounded (possible if infinite ℓ_j, u_j exist)

Step D: use Double MINOS with scaling;

1 **repeat**

Find a nonsingular basis matrix B from the columns of S so that $SP = (B\ N)$;
Find $v = P(v_B, v_N)$ with $Sv \equiv Bv_B + Nv_N = 0$;
Partition c accordingly as $c = P(c_B, c_N)$;
Solve $B^T y = c_B$;
Set $z_N \leftarrow c_N - N^T y$; $\tau \leftarrow (1 + \|y\|_\infty)\delta_2$;

until $\forall j \in N$, $z_j \le \tau$ *if* $v_j = \ell_j$, *and* $z_j \ge -\tau$ *if* $v_j = u_j$ *(optimal); or fail to find* $\ell - \delta_1 \le v \le u + \delta_1$ *(infeasible); or fail to find a new* B *(unbounded)*;

Step Q1: use Quad MINOS with scaling;
Start with the saved B from *Step D* to run loop **1** to find a new B;

Step Q2: use Quad MINOS without scaling;
Start with the saved B from *Step Q1* to run loop **1** to reach a final B;

Table 1: MINOS runtime options (defaults and those selected for each step of the DQQ procedure).

	Default	Step D	Step Q1	Step Q2
Precision	Double	Double	Quad	Quad
Scale	Yes	Yes	Yes	No
Feasibility tol δ_1	1e-6	1e-7	1e-15	1e-15
Optimality tol δ_2	1e-6	1e-7	1e-15	1e-15
Expand frequency	10000	100000	100000	100000
LU Factor tol	100.0	1.9	10.0	5.0
LU Update tol	10.0	1.9	10.0	5.0

Large ME models. COBRA can be used to stoichiometrically couple metabolic and macromolecular expression networks with single nucleotide resolution at genome-scale [24, 4]. The corresponding Metabolic and macromolecular Expression models (ME models) explicitly represent catalysis by macromolecules, and in turn, metabolites are substrates in macromolecular synthesis reactions. These reconstructions lead to the first multi-timescale and genome-scale stoichiometric models, as they account for multiple cellular functions operating on widely different timescales and typically account for about 40 percent of a prokaryote's open reading frames. A typical M model might be represented by 1000 reactions generated by hand [28]. In contrast, ME models can have more than 50,000 reactions, most of which have been generated algorithmically from template reactions (defined in the literature) and omics data [24, 4]. Typical net metabolic reaction rates are 6 orders of magnitude faster than macromolecular synthesis reaction rates (millimole/gDW vs nanomole/gDW, gDW = gram dry weight), and the number of metabolic moieties in a macromolecule can be many orders of magnitude larger than in a typical metabolite. The combined effect is that the corresponding ME models have biochemically significant digits over many orders of magnitude. When Flux Balance Analysis (FBA) is augmented with coupling constraints [29] that constrain the ratio between catalytic usage of a molecule and synthesis of the same molecule, the corresponding linear optimization problem is multiscale in the sense that both data values and solution values have greatly varying magnitudes. For a typical ME model, input data values (objective, stoichiometric or coupling coefficients, or bounds) differ by 6 orders of magnitude, and biochemically meaningful solution values can be as large as 10^8 or as small as 10^{-10}.

The results of DQQ on three large ME models TMA_ME, GlcAerWT, and GlcAlift are shown in Tables 2–3, including problem dimensions (m, n), number of nonzero entries (nnz(S)), norms of the optimal primal and dual solution vectors (v^*, y^*), number of iterations, runtime, objective value, primal and dual infeasibility after each step (Pinf and Dinf), and total solve time for each model. (The constraints in (1) are satisfied to within Pinf, and $z_j/(1 + \|y^*\|_\infty)$ has the correct sign to within Dinf, where $B^T y = c_B$ for the optimal basis B, and $z = c - S^T y$.)

TMA_ME developed by Lerman et al. [4] has some large matrix entries $|S_{ij}|$ and many small solution values v_j that are meaningful to systems biologists. For example, transcription and translation rates can have values $O(10^{-10})$ or less, which is much smaller than metabolic reactions. These small values are linked to large matrix entries arising from building large macromolecules from smaller constituents [24]. The ME part of the model also contains small matrix entries. For instance, enzyme levels are estimated in ME models by dividing certain metabolic fluxes by "effective rate constants." Because these constants are typically large (e.g., 234,000 h^{-1}), the matrix entries (the inverse of the rate constants) become small. In step D, almost all iterations went on finding a feasible solution, and the objective then had the correct order of magnitude (but only one correct digit). Step Q1 improved the accuracy, and Step Q2 provided confirmation. Note that the efficiency advantage of our approach is also evident: 385 seconds solve time for DQQ (Total time in Table 2) compared to 2 weeks using exact arithmetic [4].

Two slightly different versions of this model provided welcome empirical evidence that the optimal objective and solution values do not change significantly when the problem data are perturbed by $O(10^{-6})$ (see Supplementary Information).

GlcAerWT is a ME model from the detailed study by Thiele et al. [24]. After 33,000 iterations, Double MINOS began to report singularities following updates to the basis LU factors (71 times during the next 15,000 iterations). After 47,718 iterations (D itns in Table 2),

Table 2: Three large ME biochemical network models TMA_ME, GlcAerWT, GlcAlift [4, 24, 21]. Dimensions of $m \times n$ constraint matrices S, size of the largest optimal primal and dual variables v^*, y^*, number of iterations and runtimes in seconds for each step, and the total runtime of each model.

ME model	TMA_ME	GlcAerWT	GlcAlift
m	18210	68300	69529
n	17535	76664	77893
nnz(S)	336302	926357	928815
max $\|S_{ij}\|$	2.1e+04	8.0e+05	2.6e+05
$\|v^*\|_\infty$	5.9e+00	6.3e+07	6.3e+07
$\|y^*\|_\infty$	1.1e+00	2.4e+07	2.4e+07
D itns	21026	47718	93857
D time	350.9	10567.8	15913.7
Q1 itns	597	4287	1631
Q1 time	29.0	1958.9	277.3
Q2 itns	0	4	1
Q2 time	5.4	72.1	44.0
Total time	385	12599	16235

Table 3: Three large ME biochemical network models TMA_ME, GlcAerWT, GlcAlift [4, 24, 21]. Optimal objective value of each step, Pinf and Dinf = final maximum primal and dual infeasibilities ($\log_{10}$ values tabulated, except – means 0). Bold figures show the final *(step Q2)* Pinf and Dinf.

ME model	Step	Objective	Pinf	Dinf
TMA_ME	D	8.3789966820e−07	−06	−05
	Q1	8.7036315385e−07	−25	−32
	Q2	8.7036315385e−07	–	**−32**
GlcAerWT	D	−6.7687059922e+05	−04	+00
	Q1	−7.0382449681e+05	−07	−26
	Q2	−7.0382449681e+05	**−21**	**−22**
GlcAlift	D	−5.3319574961e+05	−03	−01
	Q1	−7.0434008750e+05	−08	−22
	Q2	−7.0434008750e+05	**−18**	**−23**

step D terminated with maximum primal and dual infeasibilities $O(10^{-4})$ and $O(1)$ (Pinf and Dinf in Table 3). These were small enough to be classified "Optimal", but we see that the final objective value −6.7687e+05 had no correct digits compared to −7.0382e+05 in steps Q1 and Q2. For large models, step Q1 is important. It required significant work: 4,287 iterations costing 1958.9 seconds (Q1 itns and time in Table 2). Step Q2 quickly confirmed the final objective value with high accuracy. This, the largest ME model so far, solved in 12,599 seconds (3.5 hours) compared to an expected time of months for an exact solver.

GlcAlift is motivated by the difficulties with solving TMA_ME and GlcAerWT in Double arithmetic. The lifting technique of [21] was applied to GlcAerWT to reduce some of the large matrix values. The aim of lifting is to remove the need for scaling (and hence the difficulties with unscaling), but with DQQ we do scale in step D because steps Q1 and Q2 follow. Our experience is that lifting improves accuracy for Double solvers but substantially increases the simplex iterations. On GlcAlift, Double MINOS again reported frequent singularities following basis updates (235 times starting near iteration 40,000). It took 93,857 iterations (D itns in Table 2), twice as many as GlcAerWT, with only a slight improvement in max{Pinf, Dinf} (Table 3). Double MINOS with scaling on the lifted model couldn't reach agreement with the final objective −7.0434008750e+05 in steps Q1 and Q2, and the total solve time increased (4.5 hours), mostly in step D. The objective function for both GlcA models is to maximize variable v_{60069}. The fact that the step D objective values have no correct digits illustrates the challenge these models present. Starting from the basis that the Double solver reaches, steps Q1 and Q2 are accurate and efficient. Theoretically, the Q2 objectives for GlcAerWt and GlcAlift should agree, but limited precision in the data files could explain why there is just 3-digit agreement.

The Tomlab interface [30] and CPLEX were used by Thiele et al. [24] to improve the results for standard Double solvers. On the NEOS server [20], Gurobi was unable to solve GlcAerWT with default parameters (numeric error after nearly 600,000 iterations). It performed considerably better on GlcAlift (about 46,000 iterations) but terminated with a warning of unscaled primal/dual residuals 1.07 and 1.22e–06. As shown above, our DQQ procedure saves researchers' effort on lifting the model, and is able to solve the original model faster (3.5 hours vs 4.5 hours).

Further tests of the DQQ procedure on challenging LO problems are reported in **Methods**. As for the ME models, the simplex method in Double MINOS usually gives a good starting point for the same simplex method in Quad MINOS. Hence, much of the work can be performed efficiently with conventional 16-digit floating-point hardware to obtain near-optimal solutions. For Quad MINOS, 34-digit floating-point operations are implemented in the compiler's Quad math library via software (on today's machines). Each simplex iteration is therefore considerably slower than by hardware, but the reward is extremely high accuracy. Of significant interest is that Quad MINOS almost invariably achieves *far more accurate solutions than requested* (see bold figures in Tables 3 and 5). This is a favorable and promising empirical finding.

Discussion

Exact solvers compute exact solutions to LO problems involving rational data. Although stoichiometric coefficients for chemical reactions are in principle integers, most genome-scale metabolic models have non-integer coefficients where the stoichiometry is known to only a few digits, e.g., a coefficient in a biomass reaction. Such a stoichiometric coefficient should not be considered exact data (to be converted into a rational number for use with an exact solver). This casts doubt on any effort to compute an exact solution for a particular FBA problem.

Exact solvers are based on rational arithmetic. There has been considerable work on their application to important problems [32, 12, 9, 4]. The use of quadruple-precision and variable-precision floating-point has also been mentioned [32, 12]. Here, we exploit Quad preci-

sion more fully on a range of larger problems, knowing that current genome-scale models will continue to grow even larger.

While today's advanced LO solvers, such as CPLEX, Gurobi, Mosek, and Xpress [5, 6, 19, 33], are effective on a wide range of large and challenging linear (and mixed integer) optimization models, the study by Thiele et al. [24] emphasizes the need for improved reliability in solving FBA and ME models in systems biology. Our DQQ procedure has demonstrated that warm starts with Quad solvers are efficient, and that the accuracy achieved exceeds requirements by a very safe margin. Kahan [18] notes that "*carrying somewhat more precision in the arithmetic than twice the precision carried in the data and available for the result will vastly reduce embarrassment due to roundoff-induced anomalies*" and that "*default evaluation in Quad is the humane option.*" The "humane" approach—use of Quad solvers—is certainly more efficient than applying exact solvers.

An intriguing question remains concerning the bold figures in Tables 3 and 5. The primal and dual solutions obtained with Quad precision are substantially more accurate than the 10^{-15} requested. The same has been true for all of the classic set of Netlib problems [34] that we have run. Kahan [18] explains that perturbations get amplified by singularities near the data. He describes a *pejorative surface* of data points where singularity exists, and expects loss of accuracy as data approaches the surface. The volume surrounding the pejorative surface is the danger zone, but: "*Arithmetic precision is usually extravagant enough if it is somewhat more than twice as [great] as the data's and the desired result's. Often that shrunken volume contains no data.*" We surmise that Kahan has anticipated our observed situation, wherein LO problems defined with double-precision data appear unlikely to be too ill-conditioned for a Quad solver.

It should be said that exact simplex solvers can also be warm-started, as noted by Gleixner et al. [16, 17]. For many models, most of the work could be done by a conventional Double solver as in our DQQ procedure, and Step Q1 could be replaced by a call to an exact solver. However, for the GlcA problems, we see in Table 2 that step Q1 performs a significant number of iterations. Thus, warm-starting an exact solver on large models could remain too expensive to be practical.

Looking ahead, we note that metabolic reconstructions of the form (1) may need to be processed before they can be treated as stoichiometrically consistent models. As discussed in [35], certain rows of S may need to be deleted according to the solution ℓ of the problem $\max \|\ell\|_0$ s.t. $S^T\ell = 0$, $\ell \geq 0$. This problem can be approximated by the linear problem

$$\begin{array}{ll} \max\limits_{z,\,\ell} & \mathbf{1}^T z \\ \text{s.t.} & S^T\ell = 0, \qquad z \leq \ell, \\ & 0 \leq z \leq \mathbf{1}\alpha, \quad 0 \leq \ell \leq \mathbf{1}\beta, \end{array} \tag{2}$$

where scalars α, β are proportional to the smallest molecular mass considered non-zero and the largest molecular mass allowed (e.g., $\alpha = 10^{-4}$, $\beta = 10^4$). Note that problem (2) involves S^T and is larger than the FBA problem (1) itself. We could not design consistent FBA models in this way unless we were sure of being able to solve (2) effectively. Our work here offers assurance of such capability.

We believe that quadruple-precision solutions are now practical for multiscale applications such as FBA and flux variability analysis (FVA) computations for ME models in systems biology [1, 36, 24, 37, 29], and that our DQQ procedure justifies increased confidence as systems biologists build ever-larger models to explore new hypotheses about metabolism and macromolecular synthesis. Our combined use of Double and Quad solvers will lead to solutions of exceptional accuracy in other areas of computational science involving multiscale optimization problems. For example, Dattorro [38] has derived an approach to analog filter design that requires a Quad linear or nonlinear solver to deal with a wide range of frequencies (which must be raised to high powers). This application, like ME models with nonlinear constraints (7), can be treated with Quad precision and binary search on a sequence of problems. We have also treated the nonlinear constraints directly with the nonlinear algorithms in Quad MINOS [11, 39].

Methods

Multiscale constraint-based modeling. Consider a network of biochemical reactions, represented by a stoichiometric matrix $S \in \mathbb{R}^{m\times n}$ with each row and column corresponding to a molecular species and biochemical reaction, respectively. S_{ij} respresents the *stoichiometry* of molecular species i participating as a substrate (negative) or product (positive) in reaction j. The evolution of molecular species concentrations with respect to time (t) is given by the ordinary differential equation

$$\frac{dx(t)}{dt} = Sv(x(t)), \tag{3}$$

where $x(t) \in \mathbb{R}^m_{\geq 0}$ is a vector of time-dependent concentrations and $v(x(t)) : \mathbb{R}^m_{\geq 0} \to \mathbb{R}^n$ is a nonlinear function of concentrations, with a form that depends on the kinetic mechanism of each reaction.

If one assumes that species concentrations are time-invariant, then the set of all steady-state reaction rates, satisfying $Sv(x) = 0$, may be approximated by the linear *steady-state constraint* $Sv = 0$, where $v \in \mathbb{R}^n$ is a vector of reaction fluxes. Thermodynamic principles and experimental data can also be used to specify lower and upper *bound constraints* on reaction fluxes $\ell \leq v \leq u$. Biochemical relationships between the rates of macromolecular synthesis and utilization can be approximated by coupling of the corresponding reaction fluxes [29], e.g., pyruvate kinase reaction flux and the synthesis flux of pyruvate kinase in a ME model [24]. Flux coupling can

be represented by bounding the ratio between two reaction fluxes with two coupling coefficients:

$$\sigma_{\min} \leq \frac{v_i}{v_j} \leq \sigma_{\max}, \tag{4}$$

where v_i and v_j are a pair of non-negative fluxes. This nonlinear constraint can be reformulated into a pair of linear *coupling constraints*

$$\sigma_{\min} v_j \leq v_i, \quad v_i \leq \sigma_{\max} v_j, \tag{5}$$

or more generally a set of linear inequalities $Cv \leq d$. In addition to the aforementioned physicochemical and biochemical contraints, one may hypothesize a biologically motivated objective. For example, in modeling a growing cell, one may hypothesize that the objective is to maximize the rate of a biomass synthesis reaction. Typically, a biomass synthesis reaction is created with experimentally determined stoichiometric coefficients, each of which represents the relative composition of a cellular biomass constituent. Optimization of a linear combination of reaction fluxes $c^T v$ leads to linear optimization problems: (1). Flux balance analysis of a ME model with coupling constraints results in an ill-scaled instance of this problem because the stoichiometric coefficients and coupling coefficients vary over many orders of magnitude.

MINOS implementation. MINOS [10, 11] is a linear and nonlinear optimization solver implemented in Fortran 77 to solve problems of the form

$$\min_v \; c^T v + \varphi(v) \;\; \text{s.t.} \;\; \ell \leq \begin{pmatrix} v \\ Sv \\ f(v) \end{pmatrix} \leq u, \tag{6}$$

where $\varphi(v)$ is a smooth nonlinear function and $f(v)$ is a vector of smooth nonlinear functions (see Supplementary Information).

Further tests of DQQ. We report results from the primal simplex solver in Double MINOS and Quad MINOS on two sets of challenging LO problems shown in Table 4. As with the M and ME models, all runs were on a 2.93 GHz Apple iMac with quad-core Intel i7, using the gfortran compiler with -O flag (GNU Fortran 5.2.0). The problems were input from files in the classical MPS format of commercial mathematical programming systems [40] with 12-character fields for all data values.

The pilot problems. These are from a set of economic models developed by Professor George Dantzig's group in the Systems Optimization Laboratory at Stanford University during the 1980s. They are available from Netlib [34] and have been used in previous computational studies (e.g., [32]). We use three examples of increasing size: pilot4, pilot, pilot87. In Table 5, three lines for each problem show the results of steps D, Q1, Q2 of the DQQ procedure.

Line 1 for pilot shows that Double MINOS with cold start and scaling (step D) required 16060 simplex iterations and 9 CPU seconds. The unscaled primal solution

Table 4: Three pilot models from Netlib [34], and eight Mészáros *problematic* problems [41]. Dimensions of $m \times n$ constraint matrices S, and size of the largest optimal primal and dual variables v^*, y^*.

model	m	n	nnz(S)	max $\lvert S_{ij}\rvert$	$\lVert v^*\rVert_\infty$	$\lVert y^*\rVert_\infty$
pilot4	411	1000	5145	2.8e+04	9.6e+04	2.7e+02
pilot	1442	3652	43220	1.5e+02	4.1e+03	2.0e+02
pilot87	2031	4883	73804	1.0e+03	2.4e+04	1.1e+01
de063155	853	1488	5405	8.3e+11	3.1e+13	6.2e+04
de063157	937	1488	5551	2.3e+18	2.3e+17	6.2e+04
de080285	937	1488	5471	9.7e+02	1.1e+02	2.6e+01
gen1	770	2560	64621	1.0e+00	3.0e+00	1.0e+00
gen2	1122	3264	84095	1.0e+00	3.3e+00	1.0e+00
gen4	1538	4297	110174	1.0e+00	3.0e+00	1.0e+00
l30	2702	15380	64790	1.8e+00	1.0e+09	4.2e+00
iprob	3002	3001	12000	9.9e+03	3.1e+02	1.1e+00

Table 5: Iterations and runtimes in seconds for step D (Double MINOS) and steps Q1, Q2 (Quad MINOS) on the problems of Table 4. Pinf and Dinf = final maximum primal and dual infeasibilities ($\log_{10}$ values tabulated, except – means 0). Problem iprob is infeasible. Bold figures show Pinf and Dinf at the end of step Q2. Note that Pinf/$\lVert v^*\rVert_\infty$ and Dinf/$\lVert y^*\rVert_\infty$ are all $O(10^{-30})$ or smaller, even though only $O(10^{-15})$ was requested. This is an unexpectedly favorable empirical finding.

model	Itns	Times	Final objective	Pinf	Dinf
pilot4	1464	0.1	−2.5811392619e+03	−05	−12
	7	0.0	−2.5811392589e+03	−52	−31
	0	0.0	−2.5811392589e+03	–	**−29**
pilot	16060	9.0	−5.5739887685e+02	−06	−03
	29	0.3	−5.5748972928e+02	–	−32
	0	0.1	−5.5748972928e+02	–	**−32**
pilot87	19340	22.6	3.0171038489e+02	−08	−06
	32	0.9	3.0171034733e+02	–	−32
	0	0.6	3.0171034733e+02	–	**−33**
de063155	973	0.1	1.8968895791e+10	−14	+03
	90	0.1	9.8830944565e+09	–	−27
	0	0.0	9.8830944565e+09	–	**−24**
de063157	1473	0.1	2.6170359397e+12	–	+08
	286	0.2	2.1528501109e+07	−29	−12
	0	0.0	2.1528501109e+07	–	**−12**
de080285	418	0.0	1.4495817688e+01	−09	−02
	132	0.1	1.3924732864e+01	−35	−32
	0	0.0	1.3924732864e+01	–	**−32**
gen1	303212	156.9	−8.1861282705e−08	−06	−13
	216746	3431.2	1.2939275026e−06	−12	−31
	8304	112.5	1.2953925804e−06	**−46**	**−31**
gen2	45905	60.0	3.2927907833e+00	−04	−12
	2192	359.9	3.2927907840e+00	–	−29
	0	10.4	3.2927907840e+00	–	**−32**
gen4	38111	151.3	−1.2724113149e−07	−07	−12
	58118	6420.2	2.8932557999e−06	−12	−31
	50	4.3	2.8933064888e−06	**−53**	**−30**
l30	1302602	805.6	9.5266141670e−01	−08	−09
	500000	6168.8	−4.5793509329e−26	−25	−00
	16292	204.4	−6.6656750251e−26	**−25**	**−31**
iprob	1087	0.2	2.6891551285e+03	+02	−11
	0	0.0	2.6891551285e+03	+02	−30
	0	0.0	2.6891551285e+03	+02	**−28**

v satisfied the constraints in (1) to within $O(10^{-6})$ and the dual solution y satisfied the optimality conditions to within $O(10^{-3})$.

Line 2 for pilot shows that Quad MINOS starting from that point with scaling (step Q1) needed only 29 iterations and 0.3 seconds to obtain a very accurate solution.

Line 3 for pilot shows that in the "insurance" step Q2, Quad MINOS warm-starting again but with no scaling gave an equally good solution (maximum infeasibilities 0.0 and $O(10^{-32})$).

The final Double and Quad objective values differ in the 4th significant digit, as suggested by removal of step D's $O(10^{-3})$ dual infeasibility.

Results for the other pilot problems are analogous.

The Mészáros problematic problems. Our DQQ procedure was initially developed for this set of difficult LO problems collected by Mészáros [41], who names them *problematic* and notes that "*modeling mistakes made these problems "crazy," but they are excellent examples to test numerical robustness of a solver.*" They were provided as MPS files [22]. The first two problems have unusually large entries in the constraint matrix S. The step D objective value for de063155 has at best 1 digit of precision, and is quite erroneous for de063157. Nevertheless, the step Q1 and Q2 solutions are seen to be highly accurate (small Pinf and Dinf values) when the solution norms are taken into account.

The gen problems come from image reconstruction, with no large entries in S, v, y but highly degenerate primal solutions v. (In steps D and Q1 for gen1, 60% of the iterations made no improvement to the objective, and the final solution has 30% of the basic variables on their lower bound.) For gen1, step Q1 gave an almost feasible initial solution (253 basic variables outside their bounds by more than 10^{-15} with a sum of infeasibilities of only $O(10^{-8})$), yet over 200,000 iterations were needed in step Q1 to reach optimality. These examples show that Quad precision does not remove the need for a more rigorous anti-degeneracy procedure (such as Wolfe's method as advocated by Fletcher [42]), and/or steepest-edge pricing [43], to reduce significantly the total number of iterations. Problems gen1 and gen4 show that step Q2 is sometimes needed to achieve high accuracy.

Problem l30 behaved similarly (80% degenerate iterations in steps D and Q1). The tiny objective value is essentially zero, so we can't expect the Q1 and Q2 objectives to agree in their leading digits. The Q1 iterations were inadvertently limited to 500,000, but step Q2 did not have much further to go.

Problem iprob is an artificial one that was intended to be feasible with a very ill-conditioned optimal basis, but the MPS file provided to us contained low-precision data (many entries like 0.604 or 0.0422). Our Double and Quad runs agree that the problem is infeasible. This is an example of Quad removing some doubt that would be inevitable with just Double.

Table 5 shows that Quad MINOS almost invariably achieves far more accurate solutions than requested, in the sense that the maximum primal and dual infeasibilities are almost always far smaller than 10^{-15}. Thus our procedure for handling the *problematic* problems is appropriate for the systems biology M and ME models. Like the gen problems, the ME models showed 40–60% degenerate iterations in step D, but fortunately not so many total iterations in step Q1 (see Table 2). This is important for FVA and for ME with nonlinear constraints, where there are many warm starts.

ME models (FBA with coupling constraints). As coupling constraints are often functions of the organism's growth rate μ, O'Brien et al. [44] consider growth-rate optimization nonlinearly with the single μ as the objective in (1) instead of via a linear biomass objective function. Nonlinear constraints of the form

$$v_i \geq \mu \sum_j v_j / k_{i,j}^{\text{eff}} \tag{7}$$

are added to (1), where v_i, v_j, μ are all variables, and $k_{i,j}^{\text{eff}}$ is an effective rate constant. Constraints (7) are linear if μ is fixed at a specific value μ_k. O'Brien et al. [44] employ a binary search on a discrete set of values within an interval $[\mu_{\min}, \mu_{\max}]$ to find the largest $\mu_k \equiv \mu^*$ that keeps the associated linear problem feasible. Thus, the procedure requires reliable solution of a sequence of related LO problems.

Flux Variability Analysis (FVA). After FBA (1) returns an optimal objective value $c^T v^* = Z_0$, FVA examines how far a particular flux v_j can vary within the feasible region without changing the optimal objective significantly (if $\gamma \approx 1$):

$$\min_v \ \pm v_j \ \text{ s.t. } \ Sv = 0, \ \ c^T v \geq \gamma Z_0, \ \ l \leq v \leq u, \tag{8}$$

where $0 < \gamma < 1$. Potentially $2n$ LO problems (8) are solved if all reactions are of interest, with warm starts being used when j is increased to $j+1$ [37].

For such a sequence of related problems, warm-starting each problem in Quad would be simplest (calling a single solver), but warm-starting in Double and then in Quad could sometimes be more efficient.

Conventional iterative refinement. A Double simplex solver would be more reliable with the help of iterative refinement (Wilkinson [18]), but we found this inadequate for the biology models (see DRR procedure in Supplementary Information).

The zoom strategy. A step toward warm-starting interior methods for optimization was proposed in [46] to take advantage of the fact that a low-accuracy solution (x_1, y_1) for a general problem

$$\min \ c^T x \ \text{ s.t. } \ Ax = b, \ \ \ell \leq x \leq u \tag{9}$$

can be obtained relatively cheaply when an iterative solver for linear systems is used to compute each search direction. (The iterative solver must work harder as the

interior method approaches a solution.) If (x_1, y_1) has at least some correct digits, the primal residual $r_1 = b - Ax_1$ will be somewhat small ($\|r_1\| = O(1/\sigma)$ for some $\sigma \gg 1$) and the dual residual $d_1 = c - A^T y_1$ will be comparably small in the elements associated with the final B. If we define

$$\begin{array}{ll} b_2 = \sigma r_1, & c_2 = \sigma d_1, \\ \ell_2 = \sigma(\ell - x_1), & u_2 = \sigma(u - x_1), \\ x = x_1 + \frac{1}{\sigma} x_2, & y = y_1 + \frac{1}{\sigma} y_2, \end{array} \tag{10}$$

and note that the problem is equivalent to

$$\min\ c^T x - y_1^T (Ax - b)\ \text{ s.t. }\ Ax = b,\ \ell \le x \le u \tag{11}$$

with dual variable $y - y_1$, we see that x_2 solves

$$\min\ c_2^T x_2\ \text{ s.t. }\ Ax_2 = b_2,\ \ell_2 \le x_2 \le u_2 \tag{12}$$

with dual variable y_2. Importantly, with σ chosen carefully we expect (x_2, y_2) in this "*zoomed in*" problem to be of order 1. Hence we can solve the problem with the same solver as before (as solvers use absolute tolerances and assume that A and the solution are of order 1). If the computed (x_2, y_2) has at least some digits of accuracy, the correction $x_1 \leftarrow x_1 + \frac{1}{\sigma} x_2$, $y_1 \leftarrow y_1 + \frac{1}{\sigma} y_2$ will be more accurate than before. The process can be repeated. With repeated zooms (named *refinement rounds* in [16, 17]), the residuals (r_1, d_1) must be computed with increasingly high precision. Subject to the expense of using rational arithmetic for this purpose, [16] gives extensive results for over 1000 challenging problems and shows that exceptional accuracy can be obtained in reasonable time: only 3 or 4 refinements to achieve 10^{-50} precision, and less than 20 refinements to achieve 10^{-250}. The SoPlex80bit solver [7, 8] is used for each refinement round with feasibility and optimality tolerances set to 10^{-9}. In [16] the authors recognize that much depends on the robustness of the simplex solver used for the original problem and each refinement. The potential difficulties are the same as in each step of our DRR procedure, where Double MINOS is on the brink of failure on the Glc problems because B is frequently near-singular when it is refactorized every 100 iterations. A practical answer for [16] is to use a more accurate floating-point solver such as Quad MINOS (or Quad versions of SoPlex or SNOPT [5]) for all refinement rounds.

DQQ serves the current purpose. In the context of ME models whose non-integer data is accurate to only 4 or 5 digits, we don't need 10^{-50} precision. Tables 3 and 5 show that our DQQ procedure achieves more accuracy than necessary on all tested examples. For models where the Double solver is expected to encounter difficulty, step D can use a reasonable iteration limit. Step Q1 will perform more of the total work with greatly improved reliability. Step Q2 provides a small but important improvement at negligible cost, ensuring small residuals for the original (unscaled) problem.

Data and software availability. Double and Quad Fortran 77 implementations of MINOS are included within the Cobra toolbox [2]. MPS or JSON files for all models discussed are available from [22]. Python code for running Double and Quad MINOS on the BiGG JSON files is also available from [22].

Acknowledgements

We thank Jon Dattorro, Philip Gill, Edward O'Brien, Elizabeth Wong, and several other colleagues for their valuable help. Joshua Lerman at UC San Diego provided the model named TMA_ME here (originally model_final_build_unscaled.mat) and advised us of the final objective values obtained by SoPlex and QSopt_ex. Yuekai Sun at Stanford University created the reformulated version named GlcAlift here. Ed Klotz of IBM in Carson City NV provided MPS files for the Mészáros *problematic* problems and lengthy discussions of their properties.

This work was supported by the National Institute of General Medical Sciences of the National Institutes of Health [award U01GM102098] and the US Department of Energy [award desc0010429]. The content is solely the responsibility of the authors and does not necessarily represent the official views of the funding agencies.

Author Contributions

R.F. and M.S. conceived this study. D.M. and M.S. developed the DQQ procedure and designed the manuscript. M.S. developed the Double and Quad MINOS solvers. L.Y. implemented Python interfaces and verified the solvers on linear and nonlinear ME models. R.F. implemented Matlab interfaces within the COBRA toolbox. I.T. highlighted the impact of coupling constraints in ME models and built the largest example, GlcAerWT. All authors read and revised the manuscript.

Competing Financial Interests

The authors declare no competing financial interests.

References

[1] B. O. Palsson. *Systems Biology: Properties of Reconstructed Networks*. Cambridge University Press, NY, 2006.

[2] J. Schellenberger, R. Que, R. M. T. Fleming, I. Thiele, J. D. Orth, A. M. Feist, D. C. Zielinski, A. Bordbar, N. E. Lewis, S. Rahmanian, et al. Quantitative prediction of cellular metabolism with constraint-based models: the COBRA Toolbox v2.0. *Nature Protocols*, 6(9):1290–1307, 2011.

[3] UCSD Systems Biology Research Group:. COBRApy, Constraints-based modeling of biological networks. `http://opencobra.github.io/cobrapy/`, 2015.

[4] Gurobi optimization system for linear and integer programming. `http://www.gurobi.com`, 2014.

[5] IBM ILOG CPLEX optimizer. `http://www.ibm.com/software/commerce/optimization/cplex-optimizer/`, 2014.

[6] MOSEK Optimization Software. `http://www.mosek.com/`, 2014.

[7] I. Thiele, N. Jamshidi, R. M. T. Fleming, and B. O. Palsson. Genome-scale reconstruction of *E. coli*'s transcriptional and translational machinery: A knowledge-base, its mathematical formulation, and its functional characterization. *PLoS Comput Biol*, 5(3):e1000312, 2009.

[8] I. Thiele, R. M. T. Fleming, A. Bordbar, R. Que, and B. O. Palsson. A systems biology approach to the evolution of codon use pattern. *(submitted, Available from Nature Precedings http://dx.doi.org/10.1038/npre.2011.6312.1)*, 2011.

[9] R. Fourer. Solving staircase linear programs by the simplex method, 1: Inversion. *Math. Program.*, 23:274–313, 1982.

[10] J. A. Tomlin. On scaling linear programming problems. In *Computational Practice in Mathematical Programming*, volume 4 of *Mathematical Programming Studies*, pages 146–166. Springer, 1975.

[11] Y. Sun, R. M. T. Fleming, I. Thiele, and M. A. Saunders. Robust flux balance analysis of multiscale biochemical reaction networks. *BMC Bioinformatics*, 14(240), 2013.

[12] D. L. Applegate, W. Cook, S. Dash, and D. G. Espinoza. Exact solutions to linear programming problems. *Operations Res. Lett.*, 35:693–699, 2007.

[13] QSopt_ex: A simplex solver for computing exact rational solutions to LP problems. `http://www.math.uwaterloo.ca/~bico/qsopt/index.html`, 2008.

[14] J. A. Lerman, D. R. Hyduke, H. Latif, V. A. Portnoy, N. E. Lewis, J. D. Orth, A. C. Schrimpe-Rutledge, R. D. Smith, J. N. Adkins, K. Zengler, and B. O. Palsson. In silico method for modelling metabolism and gene product expression at genome scale. *Nature Communications*, 3(929):10 pp., 2012.

[15] L. Chindelevitch, J. Trigg, A. Regev, and B. Berger. An exact arithmetic toolbox for a consistent and reproducible structural analysis of metabolic network models. *Nat. Comm.*, 5(4893):9 pp, 2014.

[16] A. M. Gleixner, D. E. Steffy, and K. Wolter. Iterative refinement for linear programming. ZIB Report 15-15, Konrad-Zuse-Zentrum für Informationstechnik Berlin, Berlin, Germany, May 2015.

[17] Ambros M. Gleixner. *Exact and Fast Algorithms for Mixed-Integer Nonlinear Programming*. PhD thesis, Konrad-Zuse-Zentrum für Informationstechnik Berlin (ZIB), Technical University of Berlin, 2015. `http://dx.doi.org/10.14279/depositonce-4938`.

[18] W. Kahan. Desperately needed remedies for the undebuggability of large floating-point computations in science and engineering. IFIP/SIAM/NIST Working Conference on Uncertainty Quantification in Scientific Computing, Boulder CO, `http://www.eecs.berkeley.edu/~wkahan/Boulder.pdf`, 2011.

[19] B. A. Murtagh and M. A. Saunders. Large-scale linearly constrained optimization. *Math. Program.*, 14:41–72, 1978.

[20] B. A. Murtagh and M. A. Saunders. A projected Lagrangian algorithm and its implementation for sparse nonlinear constraints. *Math. Program. Study*, 16:84–117, 1982.

[21] UCSD Systems Biology Research Group:. Available predictive genome-scale metabolic network reconstructions. `http://systemsbiology.ucsd.edu/InSilicoOrganisms/OtherOrganisms`, 2015.

[22] Multiscale systems biology collaboration. `http://stanford.edu/group/SOL/multiscale/models.html`, 2015.

[23] Z. A. King, J. S. Lu, A. Dräger, P. C. Miller, S. Federowicz, J. A. Lerman, A. Ebrahim, B. O. Palsson, and N. E. Lewis. BiGG Models: A platform for integrating, standardizing, and sharing genome-scale models. *Nucl. Acids Res.*, 2015.

[24] I. Thiele, R. M. T. Fleming, R. Que, A. Bordbar, D. Diep, and B. O. Palsson. Multiscale modeling of metabolism and macromolecular synthesis in *E. coli* and its application to the evolution of codon usage. *PLOS ONE*, 7(9):18 pp., 2012.

[25] P. E. Gill, W. Murray, M. A. Saunders, and M. H. Wright. A practical anti-cycling procedure for linear and nonlinear programming. *Math. Program.*, 45:437–474, 1989.

[26] A. Ebrahim. Generation of 83 models from UCSD repository. `https://github.com/opencobra/m_model_collection/blob/master/load_models.ipynb`, 2015.

[27] A. Ebrahim. Solution of 83 models from UCSD repository. `https://github.com/opencobra/m_model_collection/blob/master/exact_solving_models.ipynb`, 2015.

[28] A. M. Feist, C. S. Henry, J. L. Reed, M. Krummenacker, A. R. Joyce, P. D. Karp, L. J. Broadbelt, V. Hatzimanikatis, and B. O. Palsson. A genome-scale metabolic reconstruction for *Escherichia coli* K-12 MG1655 that accounts for 1260 ORFs and thermodynamic information. *Mol Sys Bio*, 3(1):e121, 2007.

[29] I. Thiele, R. M. T. Fleming, A. Bordbar, J. Schellenberger, and B. O. Palsson. Functional characterization of alternate optimal solutions of *Escherichia coli*'s transcriptional and translational machinery. *Biophysical J.*, 98(10):2072–2081, 2010.

[30] TOMLAB optimization environment for MATLAB. `http://tomopt.com`, 2015.

[31] NEOS server for optimization. `http://www.neos-server.org/neos/`.

[32] T. Koch. The final NETLIB-LP results. *Operations Research Letters*, 32:138–142, 2004.

[33] FICO Xpress Optimization Suite. `http://www.fico.com/en/products/fico-xpress-optimization-suite/`, 2015.

[34] Netlib collection of LP problems in MPS format. `http://www.netlib.org/lp/data`, 1988.

[35] R. M. T. Fleming, N. Vlassis, I. Thiele, and M. A. Saunders. Conditions for duality between fluxes and concentrations in biochemical networks. *J. Theoretical Biology,* accepted, 2016.

[36] Jeffrey D. Orth, Ines Thiele, and Bernhard O. Palsson. What is flux balance analysis? *Nature Biotechnology*, 28(3):245–248, 2010.

[37] S. Gudmundsson and I. Thiele. Computationally efficient flux variability analysis. *BMC Bioinformatics*, 11(489):3 pp., 2010.

[38] J. Dattorro. *Convex Optimization & Euclidean Distance Geometry*. Meboo Publishing USA, 2015. v2015.09.29, pp 208–214.

[39] L. Yang, D. Ma, A. Ebrahim, C. J. Lloyd, M. A. Saunders, and B. O. Palsson. solveME: fast and reliable solution of nonlinear ME models. In preparation, 2016.

[40] Input format for LP data. `http://lpsolve.sourceforge.net/5.5/mps-format.htm`, 1960.

[41] C. Mészáros. A collection of challenging LP problems. `http://www.sztaki.hu/~meszaros/public_ftp/lptestset/problematic`, 2004.

[42] R. Fletcher. On Wolfe's method for resolving degeneracy in linearly constrained optimization. *SIAM J. Optim.*, 24(3):1122–1137, 2014.

[43] J. J. Forrest and D. Goldfarb. Steepest-edge simplex algorithms for linear programming. *Math. Program.*, 57:341–374, 1992.

[44] E. J. O'Brien, J. A. Lerman, R. L. Chang, D. R. Hyduke, and B. O. Palsson. Genome-scale models of metabolism and gene expression extend and refine growth phenotype prediction. *Molecular Systems Biology*, 9(693):13 pp., 2013.

[45] J. H. Wilkinson. *The Algebraic Eigenvalue Problem*. Oxford University Press, 1965.

[46] M. A. Saunders and L. Tenenblat. The zoom strategy for accelerating and warm-starting interior methods. Talk at INFORMS Annual Meeting, Pittsburgh, PA, USA, November 2006.

[47] R. Wunderling. *Paralleler und objektorientierter Simplex-Algorithmus*. PhD thesis, Technische Universität Berlin, 1996.

[48] Soplex: The sequential object-oriented simplex solver. `http://soplex.zib.de`, 2016.

[49] P. E. Gill, W. Murray, and M. A. Saunders. SNOPT: An SQP algorithm for large-scale constrained optimization. *SIAM Review*, 47(1):99–131, 2005. SIGEST article.

Supplementary Information

DQQ procedure

Figure 1 summarizes our approach to achieving reliability and efficiency for multiscale linear and nonlinear optimization problems.

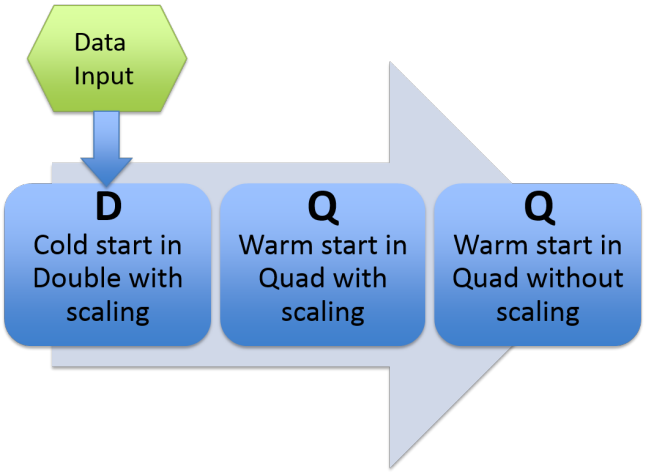

Figure 1: Flowchart for the 3-step DQQ procedure.

The main paper reports application of DQQ to three large ME models (TMA_ME, GlcAerWT, GlcAlift) and to some other challenging linear optimization problems (the pilot economic models and the Mészáros "problematic" set). Below we provide the following supplementary information:

- Solution of 78 Metabolic models by Double and Quad MINOS, verifying that the Double solver gives reliable results.
- Solution of two slightly different forms of the TMA_ME model, showing robustness of solution values with respect to $O(10^{-6})$ relative perturbations of the data.
- Some details of the Double and Quad MINOS implementations.
- Experiments with conventional iterative refinement (DRR procedure).
- Results with Gurobi on the ME models.
- Results with SoPlex80bit on the ME models.

Metabolic models with Quad solvers admit biomass synthesis

COBRA models of metabolic networks assume the existence of at least one steady-state flux vector that satisfies the imposed constraints and admits a non-zero optimal objective. Where the objective is to maximize a biomass synthesis reaction, the corresponding FBA problem should admit a nonzero biomass synthesis rate. It is established practice to solve monoscale metabolic FBA problems with Double solvers, so one may ask: do biomass synthesis predictions from metabolic models hold when higher precision solvers are applied to the same FBA problem? We tested 78 M models derived from the BiGG database [1] using Double and Quad MINOS. We downloaded these models in the JSON format and parsed them using the JSON reader in cobrapy [2]. The models were not modified after loading, so all constraints, bounds, and objective coefficients were used as in the original files. All models were feasible using both Double and Quad, and all but five models had an optimal objective value greater than zero. Of these five models, four simply had all-zero objective coefficients, while the remaining (RECON1) model maximized a single reaction (S6T14g) but its optimal value was zero. The maximum difference in objective value between Double and Quad was 2.6×10^{-12}. The additional precision provided by Quad MINOS enabled us to conclude efficiently and effectively that the 78 metabolic models could be solved reliably using a Double solver. This conclusion is consistent with previous findings by Ebrahim et al. [3].

Robustness of solution values for TMA_ME

TMA_ME [4] was the first ME model that we used for Quad experiments. The data S, c, ℓ, u came as a Matlab structure with $c_j = 0$, $\ell_j = 0$, $u_j = 1000$ for most j, except four variables had smaller upper bounds, the last variable had moderate positive bounds, and 64 variables were fixed at zero. The objective was to maximize flux v_{17533}. We output the data to a plain text file. Most entries of S were integers (represented exactly), but about 5000 S_{ij} values were of the form 8.037943687315e-01 or 3.488862338191e-06 with 13 significant digits. The text data was read into Double and Quad versions of a prototype Fortran 90 implementation of SQOPT [5].

For the present work, we used the same Matlab data to generate an MPS file for input into MINOS. Since this is limited to 6 significant digits, the values in the preceding paragraph were rounded to 8.03794e-01 and 3.48886e-06 and in total about 5000 S_{ij} values had $O(10^{-6})$ relative perturbations of this kind. This was a fortuitous limitation for the ME models. We have been concerned that such data perturbations could alter the FBA solution greatly because the final basis matrices could have condition number as large as 10^6 or even 10^{12} (as estimated by LUSOL [6] each time SQOPT or MINOS factorizes the current basis B). However, in comparing Quad SQOPT and Quad MINOS with SoPlex [7, 8] and the exact simplex solver QSopt_ex [9], we observe in Table 6 that the final objective values for TMA_ME in Matlab data reported by QSopt_ex and Quad SQOPT match in every digit. Moreover, the objective value achieved by Quad MINOS on the perturbed data in MPS format agrees to 5 digits of the results from the exact solver QSopt_ex on the "accurate" data. These results show the robustness of the TMA_ME model and our 34-digit Quad solvers.

Table 6: TMA_ME model. Robustness of objective values computed by four high-accuracy solvers for two slightly different versions of the problem with 13-digit and 6-digit data (from Matlab and MPS data respectively).

	Optimal objective	
SoPlex80bit	8.703671403e–07	Matlab data
QSopt_ex	8.703646169e–07	Matlab data
Quad SQOPT	8.703646169e–07	Matlab data
Quad MINOS	8.703631539e–07	MPS data

Table 7: TMA_ME model. Robustness of small solution values v_j and w_j computed by Quad MINOS for two slightly different versions (Matlab and MPS data respectively).

j	107	201	302
v_j	2.336815e–06	8.703646e–07	1.454536e–11
w_j	2.336823e–06	8.703632e–07	1.454540e–11

More importantly, for the most part *even small solution values* are perturbed in only the 5th or 6th significant digit. Let v and w be the solutions obtained on slightly different data. Some example values are given in Table 7. Among all j for which $\max(v_j, w_j) > \delta_1 = 10^{-15}$ (the feasibility tolerance), the largest relative difference $|v_j - w_j| / \max(v_j, w_j)$ was less than 10^{-5} for all but 31 variables. For 22 of these pairs, either v_j or w_j was primal or dual degenerate (meaning one of them was zero and there are alternative solutions with the same objective value). The remaining 9 variables had v_j, w_j values shown in Table 8.

We see that the values are small (the same magnitude as the data perturbation) but for each of the nine pairs there is about 1 digit of agreement. We could expect thousands of small solution pairs to differ more, yet for almost *all* 17535 pairs at least 5 digits agree.

These observations about two forms of TMA_ME are welcome empirical evidence of the robustness of this multiscale model. Quad solvers can be applied to evaluate the robustness of future (increasingly large) models of metabolic networks by enabling similar comparison of high-accuracy solutions for slightly different problems.

MINOS implementation

MINOS [10, 11] is a linear and nonlinear optimization solver implemented in Fortran 77 to solve problems of the form

$$\min_v \; c^T v + \varphi(v) \;\; \text{s.t.} \;\; \ell \le \begin{pmatrix} v \\ Sv \\ f(v) \end{pmatrix} \le u, \qquad (13)$$

where $\varphi(v)$ is a smooth nonlinear function and $f(v)$ is a vector of smooth nonlinear functions. The matrix S and the Jacobian of $f(v)$ are assumed to be sparse.

Let Single/Double/Quad denote the floating-point formats defined in the 2008 IEEE 754 standard [12] with about 7/16/34 digits of precision, respectively. Single is not useful in the present context, and Double may not ensure adequate accuracy for multiscale problems. This is the reason for our work. Since release 4.6 of the GCC C and Fortran compilers [13], Quad has been available via the `long double` and `real(16)` data types. Thus, we have made a Quad version of Double MINOS using the GNU gfortran compiler (GNU Fortran 5.2.0).

Table 8: TMA_ME model. The values of 9 fluxes v_j, w_j computed by Quad MINOS for two slightly different versions of the problem, revealing robustness of all 9 solution pairs. These values have 1 digit of agreement. Almost all 17535 pairs of values agree to 5 or more digits.

j	v_j	w_j	Relative difference
16383	6.07e–07	2.04e–06	0.70
16459	1.71e–06	2.18e–06	0.22
16483	2.47e–06	5.99e–07	0.76
16730	1.44e–06	7.87e–07	0.46
17461	1.71e–06	2.18e–06	0.22
17462	2.47e–06	5.99e–07	0.76
17478	6.07e–07	2.04e–06	0.70
17507	1.44e–06	7.87e–07	0.46
17517	8.70e–07	2.97e–06	0.71

On today's machines, Double is implemented in hardware, while Quad (if available) is typically implemented in a software library, in this case GCC libquadmath [14].

For Double MINOS, floating-point variables are declared `real(8)` ($\approx$ 16 digits). For Quad MINOS, they are `real(16)` ($\approx$ 34 digits) with the data S, c, ℓ, u stored in Quad even though they are not known to that precision. This allows operations such as Sv and $S^T y$ to be carried out directly on the elements of S and the Quad vectors v, y. If S were stored in Double, such products would require each entry S_{ij} to be converted from Double to Quad at runtime (many times).

The primal simplex solver in MINOS includes geometric mean scaling [15], the EXPAND anti-degeneracy procedure [16], and partial pricing (but no steepest-edge pricing, which would generally reduce total iterations and time). Basis LU factorizations and updates are handled by LUSOL [6]. Cold starts use a Crash procedure to find a triangular initial basis matrix. Basis files are used to preserve solutions between runs and to enable warm starts.

Scaling is commonly applied to linear programs to make the scaled data and solution values closer to 1. Feasibility and optimality tolerances can be chosen more easily for the scaled problem, and LU factors of the basis matrix are more likely to be sparse. For geometric mean scaling, several passes are made through the columns and rows of S to compute a scale factor for each column and row. A difficulty is that the scaled problem may solve to within specified feasibility and optimality tolerances, but when the solution is unscaled it may lie significantly outside the original (unscaled) bounds.

EXPAND tries to accommodate consecutive "degener-

ate" simplex iterations that make no improvement to the objective function. The problem bounds are effectively expanded a tiny amount each iteration to permit nonzero improvement. Convergence is usually achieved but is not theoretically guaranteed [17]. Progress sometimes stalls for long sequences of iterations.

LUSOL bounds the subdiagonals of L when the current basis matrix B is factorized as $P_1BP_2 = LU$ with some permutations P_1, P_2. It also bounds off-diagonal elements of elementary triangular factors L_j that update L in product form each simplex iteration. (The diagonals of L and each L_j are implicitly 1.) Maximum numerical stability would be achieved by setting the LU Factor and Update tolerances to be near 1.0, but larger values are typically chosen to balance stability with sparsity. For safety, we specify 1.9 in step D of DQQ. This value guards against unstable factorization of the deceptive matrix diag(−1 2 1), and improves the reliability of Double MINOS in the present context.

Conventional iterative refinement

For the biology models, our aim is to satisfy Feasibility and Optimality tolerances of 10^{-15} (close to Double precision). It is reasonable to suppose that this could be achieved within a Double simplex solver by implementing iterative refinement (Wilkinson [18]) for every linear system involving the basis matrix B or B^T. This is a more sparing use of Quad precision than our DQQ procedure. For example, each time the current B is factorized directly (typically a new sparse LU factorization every 100 iterations), the constraints $Sv = 0$ can be satisfied more accurately by computing the primal residual $r = 0 - Sv$ from the current solution v, solving $B\Delta v_B = r$, and updating $v_B \leftarrow v_B + \Delta v_B$. In general, the new v will not be significantly more accurate unless r is computed in Quad. (If B is nearly singular, more than one refinement may be needed.) Similarly for solving $B^Ty = c_B$ after refactorization, and for two systems of the form $Bp = a$ and $B^Ty = c_B$ each iteration of the simplex method.

By analogy with DQQ, we implemented the following procedure within a test version of Double MINOS. Note that "iterative refinement" in steps R1, R2 means a single refinement for each B or B^T system, with residuals $-Sv$, $a - Bp$, $c_B - B^Ty$ computed in Quad as just described.

DRR procedure.

Step D (Cold start with scaling): Apply Double MINOS with moderately strict options. Save the final basis.

Step R1 (Warm start with refinement and scaling): Start Double MINOS from the saved basis with stricter tolerances and iterative refinement. Save the final basis.

Step R2 (Warm start with refinement but no scaling): Start Double MINOS from the second saved basis without scaling but with stricter LU tolerances and iterative refinement.

Step D is the same as for DQQ (with no refinement). The runtime options for each step are the same as for DQQ, except in steps R1, R2 the tolerances 1e–15 were relaxed to 1e–9.

Table 9: DRR procedure on three ME models. Iterations and runtimes in seconds for step D (Double MINOS with scaling) and steps R1, R2 (Double MINOS with iterative refinement, with and without scaling). Pinf and Dinf = final maximum primal and dual infeasibilities ($\log_{10}$ values tabulated). Bold figures show Pinf and Dinf at the end of step R2. The fourth line for each model shows the correct objective value (from step Q2 of DQQ).

model	Itns	Times	Final objective	Pinf	Dinf
TMA_ME	21026	350.9	8.3789966820e−07	−06	−05
	422	25.4	8.6990918717e−07	−08	−07
	71	0.0	8.7035701805e−07	**−10**	**−10**
			8.7036315385e−07		
GlcAerWT	47718	10567.8	−6.7687059922e+05	−04	+00
	907	1442.7	−7.0344344753e+05	−04	−04
	157	151.2	−7.0344342883e+05	**−10**	**−02**
			−7.0382449681e+05		
GlcAlift	19340	15913.7	−5.3319574961e+05	−03	−01
	447	198.8	−7.0331052509e+05	−03	−03
	460	0.6	−7.0330602383e+05	**−06**	**−10**
			−7.0434008750e+05		

In Table 9 we see that this simplified (cheap) form of iterative refinement is only partially successful, with step R2 achieving only 4, 3, and 2 correct digits in the final objective. For GlcAerWT, steps R1 and R2 encountered frequent near-singularities in the LU factors of B (requiring excessive refactorizations and alteration of B), and in step R2, the single refinement could not always achieve full Double precision accuracy for each system. Additional refinements would improve the final Pinf and Dinf, but would not reduce the excessive factorizations. We conclude that on the bigger ME problems, a Double solver is on the brink of failure even with the aid of conventional (Wilkinson-type) iterative refinement of each system involving B and B^T. We conclude that our DQQ procedure is a more expensive but vitally more robust approach.

Results with NEOS/Gurobi

For large linear models, commercial solvers have reached a high peak of efficiency. It would be ideal to make use of them to the extent possible. For example, their Presolve capability allows most of the optimization to be performed on a greatly reduced form of any typical model.

Table 10 summarizes the performance of Gurobi [19] on three large ME models via the NEOS server [20]. The first three results used Gurobi's default runtime options, including Presolve, Dual simplex, and Scaling (with default FeasibilityTol = OptimalityTol = 1e–6). TMA_ME seemed to solve successfully, but from the Quad MINOS solution we know that Gurobi's final objective value has no correct digits. GlcAerWT failed with "Numeric error" after many expensive iterations using 80-bit floating-point. GlcAlift also switched to 80-bit floating-point. The scaled problem seemed to solve successfully, but un-

Table 10: Performance of Gurobi with default options on three ME models. Note that "switch to quad" means switch to 80-bit floating-point (not to IEEE Quad precision). This did not help GLcAerWT. For GlcAlift2, the options were NumericFocus 3, no Presolve, and no scaling.

TMA_ME	Presolve	$18209 \times 17535 \rightarrow 2386 \times 2925$
Optimal	Iterations	1703
0.5 secs	Objective	9.6318438361E-07
	True obj	8.7036315385E-07
GlcAerWT	Presolve	$68299 \times 76664 \rightarrow 18065 \times 26157$
	Warning	switch to quad (itns $\approx$ 14000)
Numeric error	Iterations	593819
3715 secs	Objective	3.2926249E+07
	True obj	-7.0382449681E+05
GlcAlift	Presolve	$69528 \times 77893 \rightarrow 18063 \times 26155$
	Warning	switch to quad (itns $\approx$ 10000)
Optimal	Iterations	45947
109 secs	Objective	-7.043390954E+05
	True obj	-7.0434008750E+05
	Warning	unscaled primal/dual residuals: 1.07, 1.22E-06
GlcAlift2		
Optimal	Iterations	128596
844 secs	Objective	-7.043415774E+05
	True obj	-7.0434008750E+05
	Warning	unscaled primal residual: 1.05E-05

Table 11: Performance of SoPlex80bit with default options on three ME models.

TMA_ME	Simplifier	$18209 \times 17535 \rightarrow 10865 \times 11272$
Optimal	Iterations	12292
15 secs	Objective	7.49100071E-07
	True obj	8.7036315385E-07
GlcAerWT	Simplifier	$68299 \times 76664 \rightarrow 63897 \times 71008$
Optimal	Iterations	74526
765 secs	Objective	-7.03824497E+05
	True obj	-7.0382449681E+05
GlcAlift		69528×77893
Infeasible	Iterations	95203
1010 secs	Objective	1.40859732E+11
	True obj	-7.0434008750E+05

scaling damaged the primal residual and this casts significant doubt on the final solution. (This is the reason for our research.)

For GlcAlift2 we specified NumericFocus 3 with no Presolve and no scaling. These options are appropriate for lifted models [21]. Gurobi did not switch to 80-bit arithmetic, yet achieved 5 correct digits in the objective. This helps confirm the value of the lifting strategy of [21], and would provide a good starting point for steps Q1 and Q2 of DQQ. However, DQQ permits us to solve the original model GlcAerWT directly (without the lifting transformation).

Results with NEOS/SoPlex80bit

Table 11 summarizes the performance of SoPlex80bit [8] on the three ME models via NEOS with default options.

SoPlex80bit was very efficient on all ME models. For TMA_ME the final objective value has the right order of magnitude but 0 correct digits (1 less than Double MINOS). For GlcAerWT, the most difficult case, the final objective is correct in all 9 digits reported. Note that we could not judge this admirable performance without our results from DQQ.

The solver's Simplifier reduced the size of the first two models significantly. Fortuitously it was not activated on GlcAlift, as if SoPlex knew that Presolve should not be used with lifted models. However, the final status and objective value ("problem is solved [infeasible]" and 1.4e+11) were incorrect. This result (like Gurobi on GlcAerWT) emphasizes the importance of DQQ achieving 20 or more correct digits in all cases.

Looking ahead

The large-scale optimizer SNOPT [5] is maintained as a Fortran 77 solver `snopt7` [22] suitable for step D of the DQQ procedure. An accompanying Fortran 2003 version `snopt9` has also been developed, for which Double and Quad libraries can be built with only one line of source code changed. They are ideal for applying DQQ to future multiscale linear and nonlinear optimization models in numerous fields.

References

[1] Z. A. King, J. S. Lu, A. Dräger, P. C. Miller, S. Federowicz, J. A. Lerman, A. Ebrahim, B. O. Palsson, and N. E. Lewis. BiGG Models: A platform for integrating, standardizing, and sharing genome-scale models. *Nucl. Acids Res.*, 2015.

[2] Ali Ebrahim, Joshua A Lerman, Bernhard O Palsson, and Daniel R Hyduke. COBRApy: COnstraints-Based Reconstruction and Analysis for Python. *BMC Syst Biol*, 7:74, 2013.

[3] Ali Ebrahim, Eivind Almaas, Eugen Bauer, Aarash Bordbar, Anthony P Burgard, Roger L Chang, Andreas Dräger, Iman Famili, Adam M Feist, Ronan MT Fleming, et al. Do genome-scale models need exact solvers or clearer standards? *Molecular Systems Biology*, 11:831, 2015.

[4] J. A. Lerman, D. R. Hyduke, H. Latif, V. A. Portnoy, N. E. Lewis, J. D. Orth, A. C. Schrimpe-Rutledge, R. D. Smith, J. N. Adkins, K. Zengler, and B. O. Palsson. In silico method for modelling metabolism and gene product expression at genome scale. *Nature Communications*, 3(929):10 pp., 2012.

[5] P. E. Gill, W. Murray, and M. A. Saunders. SNOPT: An SQP algorithm for large-scale constrained optimization. *SIAM Review*, 47(1):99–131, 2005. SIGEST article.

[6] LUSOL sparse LU factorization package. `http://stanford.edu/group/SOL/software/lusol`, 2013.

[7] R. Wunderling. *Paralleler und objektorientierter Simplex-Algorithmus*. PhD thesis, Technische Universität Berlin, 1996.

[8] Soplex: The sequential object-oriented simplex solver. `http://soplex.zib.de`, 2016.

[9] QSopt_ex: A simplex solver for computing exact rational solutions to LP problems. http://www.math.uwaterloo.ca/~bico/qsopt/index.html, 2008.

[10] B. A. Murtagh and M. A. Saunders. Large-scale linearly constrained optimization. *Math. Program.*, 14:41–72, 1978.

[11] B. A. Murtagh and M. A. Saunders. A projected Lagrangian algorithm and its implementation for sparse nonlinear constraints. *Math. Program. Study*, 16:84–117, 1982.

[12] IEEE Std 754-2008. IEEE Computer Society, 2008.

[13] GCC, the GNU Compiler Collection. https://gcc.gnu.org/, 2016.

[14] The GCC Quad-precision math library application programming interface (API). http://gcc.gnu.org/onlinedocs/libquadmath/, 2016.

[15] R. Fourer. Solving staircase linear programs by the simplex method, 1: Inversion. *Math. Program.*, 23:274–313, 1982.

[16] P. E. Gill, W. Murray, M. A. Saunders, and M. H. Wright. A practical anti-cycling procedure for linear and nonlinear programming. *Math. Program.*, 45:437–474, 1989.

[17] J. A. J. Hall and K. I. M. McKinnon. The simplest examples where the simplex method cycles and conditions where EXPAND fails to prevent cycling. *Math. Progam., Ser. B*, 100:133–150, 2004.

[18] J. H. Wilkinson. *The Algebraic Eigenvalue Problem.* Oxford University Press, 1965.

[19] Gurobi optimization system for linear and integer programming. http://www.gurobi.com, 2014.

[20] NEOS server for optimization. http://www.neos-server.org/neos/.

[21] Y. Sun, R. M. T. Fleming, I. Thiele, and M. A. Saunders. Robust flux balance analysis of multiscale biochemical reaction networks. *BMC Bioinformatics*, 14(240), 2013.

[22] UCSD/Stanford optimization software. http://ccom.ucsd.edu/~optimizers/, 2016.